

AMENDMENTS TO THE CLAIMS

Claims 1-47 (Canceled)

1 48. (New) A virtual computer system, the virtual computer system comprising:
2 a physical computer system with a first physical processor and a second
3 physical processor;
4 a virtualization software running on the physical computer system, the
5 virtualization software supporting a virtual machine with a first virtual processor
6 and a second virtual processor, the virtual machine running a guest operating
7 system (OS) with a guest OS page table, the virtualization software generating a
8 first shadow page table and a second shadow page table derived from entries in
9 the guest OS page table, the first shadow page table being used by the first
10 physical processor and the second shadow page table being used by the second
11 physical processor, the guest OS page table containing a first primary entry
12 providing an address translation for a first virtual address and the first shadow
13 page table containing a first derived entry that was derived from the first primary
14 entry, the first derived entry also providing an address translation for the first
15 virtual address, the guest OS page table also containing a second primary entry
16 providing an address translation for a second virtual address and the second
17 shadow page table containing a second derived entry that was derived from the
18 second primary entry, the second derived entry also providing an address
19 translation for the second virtual address;
20 a producer associated with the first virtual processor;
21 a first consumer associated with the first virtual processor;
22 a second consumer associated with the second virtual processor; and
23 a patch channel matrix, the patch channel matrix comprising a first patch
24 channel by which the producer communicates information to the first consumer
25 and a second patch channel by which the producer communicates information to
26 the second consumer,

27 wherein the producer responds to a first modification to the first primary
28 entry, the first modification causing a first incoherency between the first primary
29 entry and the first derived entry, the producer writing a first information to the first
30 patch channel indicating that a change has been made to a translation in the
31 guest OS page table for the first virtual address, the first virtual processor
32 executing while the first incoherency exists,

33 wherein the producer responds to a second modification to the second
34 primary entry, the second modification causing a second incoherency between
35 the second primary entry and the second derived entry, the producer writing a
36 second information to the second patch channel indicating that a change has
37 been made to a translation in the guest OS page table for the second virtual
38 address, the second virtual processor executing while the second incoherency
39 exists,

40 wherein the first consumer reads the first information from the first patch
41 channel and, on occurrence of a first coherency event at which the first shadow
42 page table is not to be incoherent with the guest OS page table with respect to
43 the first virtual address, the first consumer eliminates the first incoherency, and

44 wherein the second consumer reads the second information from the
45 second patch channel and, on occurrence of a second coherency event at which
46 the second shadow page table is not to be incoherent with the guest OS page
47 table with respect to the second virtual address, the second consumer eliminates
48 the second incoherency.

1 49. (New) The virtual computer system of claim 48, wherein the producer also
2 provides to the first consumer a third information indicating a derived value based on
3 the first modification and a fourth information that can be used to determine whether the
4 third information remains valid, and the first consumer uses the third information to
5 modify the first derived entry to make the first derived entry coherent with the first
6 primary entry if the fourth information indicates that the third information is valid.

1 50. (New) The virtual computer system of claim 49, wherein, on occurrence
2 of the first coherency event, if the fourth information indicates that the third information
3 is no longer valid, the first consumer eliminates the first incoherency by invalidating the
4 first derived entry.

1 51. (New) The virtual computer system of claim 49, wherein, on occurrence
2 of the first coherency event, if the fourth information indicates that the third information
3 is no longer valid, the first consumer determines a current value for the first primary
4 entry and uses this current value to modify the first derived entry to make the first
5 derived entry coherent with the first primary entry.

1 52. (New) The virtual computer system of claim 49, wherein the third
2 information and the fourth information comprise distinct data values.

1 53. (New) The virtual computer system of claim 48, wherein the first
2 coherency event comprises encountering a computer instruction that invalidates an
3 entry within a Translation Lookaside Buffer.

1 54. (New) The virtual computer system of claim 48, wherein the first
2 coherency event comprises a situation in which a Translation Lookaside Buffer is
3 flushed.

1 55. (New) The virtual computer system of claim 48, wherein the first patch
2 channel is used exclusively by the producer and the first consumer, and wherein the
3 producer can write to the first patch channel at the same time that the first consumer
4 reads from the patch channel, without any data races.

1 56. (New) The virtual computer system of claim 55, wherein an entry in the
2 first patch channel is indicated as being in-use by the producer at a point in time that is
3 no earlier than the time when the rest of the entry is written to the first patch channel,

4 and wherein the entry in the first patch channel is indicated as no longer being in-use by
5 the first consumer at a point in time that is no earlier than the time when the rest of the
6 entry is read from the first patch channel.

1 57. (New) The virtual computer system of claim 48, wherein the first patch
2 channel comprises a half-duplex communication channel that is used exclusively by the
3 producer and the first consumer.

1 58. (New) The virtual computer system of claim 57, wherein the first patch
2 channel also comprises a set data structure stored in memory that is shared by the
3 producer and the first consumer.

1 59. (New) The virtual computer system of claim 58, wherein the set data
2 structure comprises a data array containing a plurality of entries.

1 60. (New) The virtual computer system of claim 48, wherein the first
2 modification to the first primary entry triggers a write trace that has been placed on the
3 guest OS page table, and the producer writes the first information to the first patch
4 channel in response to the triggering of the write trace.

1 61. (New) The virtual computer system of claim 48, wherein the first virtual
2 address and the second virtual address are the same address and the first primary
3 entry and the second primary entry are the same entry.

1 62. (New) The virtual computer system of claim 48, wherein a single shadow
2 page table comprises both the first shadow page table and the second shadow page
3 table.

1 63. (New) A method for maintaining coherency between address translations
2 in a guest OS page table and address translations in shadow page tables in a virtual
3 computer system, the virtual computer system comprising:
4 a physical computer system with a first physical processor and a second
5 physical processor;
6 a virtualization software running on the physical computer system, the
7 virtualization software supporting a virtual machine with a first virtual processor
8 and a second virtual processor, the virtual machine running a guest operating
9 system (OS) with a guest OS page table, the virtualization software generating a
10 first shadow page table and a second shadow page table derived from entries in
11 the guest OS page table, the first shadow page table being used by the first
12 physical processor and the second shadow page table being used by the second
13 physical processor, the guest OS page table containing a first primary entry
14 providing an address translation for a first virtual address and the first shadow
15 page table containing a first derived entry that was derived from the first primary
16 entry, the first derived entry also providing an address translation for the first
17 virtual address, the guest OS page table also containing a second primary entry
18 providing an address translation for a second virtual address and the second
19 shadow page table containing a second derived entry that was derived from the
20 second primary entry, the second derived entry also providing an address
21 translation for the second virtual address;
22 a producer associated with the first virtual processor;
23 a first consumer associated with the first virtual processor;
24 a second consumer associated with the second virtual processor; and
25 a patch channel matrix, the patch channel matrix comprising a first patch
26 channel by which the producer communicates information to the first consumer
27 and a second patch channel by which the producer communicates information to
28 the second consumer,
29 the method comprising:

30 responding by the producer to a first modification to the first primary entry,
31 the first modification causing a first incoherency between the first primary entry
32 and the first derived entry, and writing by the producer a first information to the
33 first patch channel indicating that a change has been made to a translation in the
34 guest OS page table for the first virtual address, the first virtual processor
35 executing while the first incoherency exists;

36 responding by the producer to a second modification to the second
37 primary entry, the second modification causing a second incoherency between
38 the second primary entry and the second derived entry, and writing by the
39 producer a second information to the second patch channel indicating that a
40 change has been made to a translation in the guest OS page table for the
41 second virtual address, the second virtual processor executing while the second
42 incoherency exists;

43 reading by the first consumer the first information from the first patch
44 channel and, on occurrence of a first coherency event at which the first shadow
45 page table is not to be incoherent with the guest OS page table with respect to
46 the first virtual address, eliminating by the first consumer the first incoherency;
47 and

48 reading by the second consumer the second information from the second
49 patch channel and, on occurrence of a second coherency event at which the
50 second shadow page table is not to be incoherent with the guest OS page table
51 with respect to the second virtual address, eliminating by the second consumer
52 the second incoherency.

1 64. (New) The method of claim 63, further comprising providing by the
2 producer to the first consumer a third information indicating a derived value based on
3 the first modification and a fourth information that can be used to determine whether the
4 third information remains valid, and using by the first consumer the third information to
5 modify the first derived entry to make the first derived entry coherent with the first
6 primary entry if the fourth information indicates that the third information is valid.

1 65. (New) The method of claim 63, wherein the first patch channel is used
2 exclusively by the producer and the first consumer, and wherein the producer can write
3 to the first patch channel at the same time that the first consumer reads from the patch
4 channel, without any data races.

1 66. (New) The method of claim 65, wherein an entry in the first patch channel
2 is indicated as being in-use by the producer at a point in time that is no earlier than the
3 time when the rest of the entry is written to the first patch channel, and wherein the
4 entry in the first patch channel is indicated as no longer being in-use by the first
5 consumer at a point in time that is no earlier than the time when the rest of the entry is
6 read from the first patch channel.

1 67. (New) The method of claim 63, wherein the first virtual address and the
2 second virtual address are the same address and the first primary entry and the second
3 primary entry are the same entry.

1 68. (New) The method of claim 63, wherein a single shadow page table
2 comprises both the first shadow page table and the second shadow page table.

1 69. (New) A computer program embodied in a computer readable medium,
2 the computer program being executable in a virtual computer system, the virtual
3 computer system comprising:

4 a physical computer system with a first physical processor and a second
5 physical processor; and

6 a virtualization software running on the physical computer system, the
7 virtualization software supporting a virtual machine with a first virtual processor
8 and a second virtual processor, the virtual machine running a guest operating
9 system (OS) with a guest OS page table, the virtualization software generating a
10 first shadow page table and a second shadow page table derived from entries in
11 the guest OS page table, the first shadow page table being used by the first
12 physical processor and the second shadow page table being used by the second
13 physical processor, the guest OS page table containing a first primary entry
14 providing an address translation for a first virtual address and the first shadow
15 page table containing a first derived entry that was derived from the first primary
16 entry, the first derived entry also providing an address translation for the first
17 virtual address, the guest OS page table also containing a second primary entry
18 providing an address translation for a second virtual address and the second
19 shadow page table containing a second derived entry that was derived from the
20 second primary entry, the second derived entry also providing an address
21 translation for the second virtual address,

22 the computer program comprising:

23 a producer associated with the first virtual processor;

24 a first consumer associated with the first virtual processor; and

25 a second consumer associated with the second virtual processor,

26 wherein the producer, the first consumer and the second consumer use a
27 patch channel matrix, the patch channel matrix comprising a first patch channel
28 by which the producer communicates information to the first consumer and a
29 second patch channel by which the producer communicates information to the
30 second consumer,

31 wherein the producer responds to a first modification to the first primary
32 entry, the first modification causing a first incoherency between the first primary
33 entry and the first derived entry, the producer writing a first information to the first
34 patch channel indicating that a change has been made to a translation in the
35 guest OS page table for the first virtual address, the first virtual processor
36 executing while the first incoherency exists,

37 wherein the producer responds to a second modification to the second
38 primary entry, the second modification causing a second incoherency between
39 the second primary entry and the second derived entry, the producer writing a
40 second information to the second patch channel indicating that a change has
41 been made to a translation in the guest OS page table for the second virtual
42 address, the second virtual processor executing while the second incoherency
43 exists,

44 wherein the first consumer reads the first information from the first patch
45 channel and, on occurrence of a first coherency event at which the first shadow
46 page table is not to be incoherent with the guest OS page table with respect to
47 the first virtual address, the first consumer eliminates the first incoherency, and

48 wherein the second consumer reads the second information from the
49 second patch channel and, on occurrence of a second coherency event at which
50 the second shadow page table is not to be incoherent with the guest OS page
51 table with respect to the second virtual address, the second consumer eliminates
52 the second incoherency.

1 70. (New) The computer program of claim 69, wherein the producer also
2 provides to the first consumer a third information indicating a derived value based on
3 the first modification and a fourth information that can be used to determine whether the
4 third information remains valid, and the first consumer uses the third information to
5 modify the first derived entry to make the first derived entry coherent with the first
6 primary entry if the fourth information indicates that the third information is valid.

1 71. (New) The computer program of claim 69, wherein the first patch channel
2 is used exclusively by the producer and the first consumer, and wherein the producer
3 can write to the first patch channel at the same time that the first consumer reads from
4 the patch channel, without any data races.

1 72. (New) The computer program of claim 71, wherein an entry in the first
2 patch channel is indicated as being in-use by the producer at a point in time that is no
3 earlier than the time when the rest of the entry is written to the first patch channel, and
4 wherein the entry in the first patch channel is indicated as no longer being in-use by the
5 first consumer at a point in time that is no earlier than the time when the rest of the entry
6 is read from the first patch channel.

1 73. (New) The computer program of claim 69, wherein the first virtual address
2 and the second virtual address are the same address and the first primary entry and the
3 second primary entry are the same entry.

1 74. (New) The computer program of claim 69, wherein a single shadow page
2 table comprises both the first shadow page table and the second shadow page table.